



All contexts are not created equal: Social stimuli win the competition for organizing reinforcement learning in 9-month-old infants

Denise M. Werchan¹ | Dima Amso²

¹Department of Population Health, New York University School of Medicine, New York, NY, USA

²Department of Psychology, Columbia University, New York, NY, USA

Correspondence

Denise Werchan, Department of Population Health, New York University School of Medicine, 227 E 30th St, 7th Fl, New York, NY 10016, USA.

Funding information

National Institutes of Health, Grant/Award Number: R21 MH113870; NSF, Grant/Award Number: 2051819

Abstract

Previous work has shown that infants as young as 8 months of age can use certain features of the environment, such as the shape or color of visual stimuli, as cues to organize simple inputs into hierarchical rule structures, a robust form of reinforcement learning that supports generalization of prior learning to new contexts. However, especially in cluttered naturalistic environments, there are an abundance of potential cues that can be used to structure learning into hierarchical rule structures. It is unclear how infants determine what features constitute a higher-order context to organize inputs into hierarchical rule structures. Here, we examine whether 9-month-old infants are biased to use social stimuli, relative to non-social stimuli, as a higher-order context to organize learning of simple visuospatial inputs into hierarchical rule sets. Infants were presented with four face/color-target location pairings, which could be learned most simply as individual associations. Alternatively, infants could use the faces or colorful backgrounds as a higher-order context to organize the inputs into simpler color-location or face-location rules, respectively. Infants were then given a generalization test designed to probe *how* they learned the initial pairings. The results indicated that infants appeared to use the faces as a higher-order context to organize simpler color-location rules, which then supported generalization of learning to new face contexts. These findings provide new evidence that infants are biased to organize reinforcement learning around social stimuli.

KEYWORDS

infants, reinforcement learning, rule learning, social cognition

1 | INTRODUCTION

The ability to structure learning and action in ways that promote the procurement of favorable outcomes is essential for adaptive behavior across the lifespan. To support this ability, a learner can use relevant cues to structure learning into flexible representations that prevent interference and support generalization of prior knowledge. For instance, an infant raised in a bilingual environment may receive

conflicting inputs when learning new information (e.g., sometimes a red round object is called “apple” and other times it is called “manzana”). To learn these conflicting inputs, the infant may use cues, such as the speaker's identity, as a higher-order context to dissociate and organize inputs into flexible rule structures (Figure 1a). Yet, it is unclear how infants determine *what* stimuli constitute a higher-order context to guide learning. Here we examine whether infants are biased to use social stimuli to organize inputs into hierarchical rule sets.

Previous work shows that humans spontaneously organize arbitrary inputs into hierarchical rule sets, even when these inputs have no inherent structure (Collins et al., 2014; Collins & Frank, 2013). This is accomplished by using certain features as higher-order contexts to organize the inputs into hierarchically organized rule sets (Figure 1b). For example, when 8- and 9-month-old infants are presented with simple cues varying by shape and color in a visuospatial learning task, they are able to use specific stimulus features, such as the shape of the cues, as a higher-order context to organize inputs into simpler color-location rules (e.g., if the context is “square,” then “red” predicts a cartoon in “location 1” and “blue” predicts a cartoon in “location 2”; Werchan & Amso, 2020b; Werchan et al., 2015). Importantly, organizing inputs in this way supports generalization to new contexts (e.g., to new shapes). Related work shows that 8- to 9-month-old infants can also use social information to organize object-label mappings and event sequences into dissociable rule sets (Werchan & Amso, 2020a; Werchan et al., 2016). However,

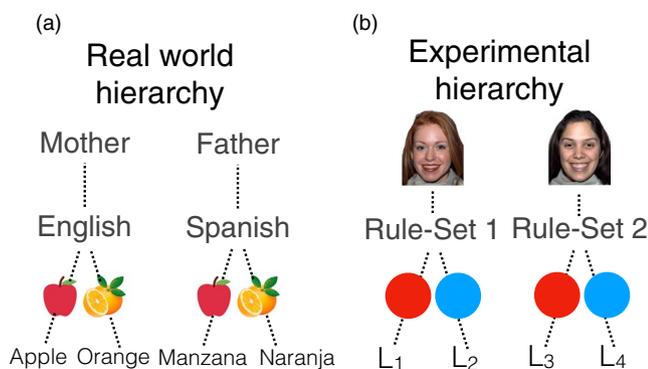


FIGURE 1 Examples of a hierarchical structure in a real-world environment, where individuals serve as contexts to organize conflicting labels for objects (a) and in an experimental environment, where faces serve as contexts to organize color-location rules (b).

especially in cluttered naturalistic environments, there are an abundance of cues that can be used to structure learning. This raises the question of how an infant learner determines what cues to use to structure inputs into flexible rule structures?

Insight into this question may be gleaned from computational models of structure learning. In the computational literature, higher-order contexts are learned through reinforcement, where a hypothesized context is positively reinforced through dopaminergic reward signals when it correctly cues an appropriate rule set (Collins & Frank, 2013). Over time, agents can use these reward histories to determine which cues signify higher-order contexts that can be used to flexibly update rule sets into working memory. As such, one possibility is that infants may be biased to use stimuli that are intrinsically rewarding as a higher-order context to organize inputs into hierarchical rule structures.

In infancy, social stimuli such as caregivers are typically the main source of biological and emotional rewards, such

Research Highlights

- Examined whether infants are biased to use social stimuli, over non-social stimuli, to organize reinforcement learning (RL) using an incidental hierarchical rule learning task.
- Presented infants with visuospatial inputs that could be learned using social or non-social stimuli as a higher-order context to organize inputs into hierarchical rule sets.
- Results indicated that infants used social stimuli to organize inputs into hierarchical rule sets, which supported generalization or prior learning to new contexts.
- This work contributes new evidence that infants are biased to organize RL around social stimuli.

as food, warmth, and comfort. Indeed, prior work has found that 7-month-old infants have increased pupil dilation and eye blink rate, which are purported physiological indices of dopaminergic reward activity, in response to the primary caregiver, which in turn drives learning of spatiotemporal patterns (Tummeltshammer et al., 2019). At the behavioral level, there is ample evidence that infants use social information to guide attention and learning from early infancy (Csibra & Gergely, 2006; Hood et al., 1998; Meltzoff et al., 2009; Scaife & Bruner, 1975). For instance, from early postnatal life, infants show a preference for infant-directed over adult-directed speech (Cooper & Aslin, 1990) and to attend to faces over non-face patterns (Johnson, 2005). Infants also follow gaze (Farroni et al., 2004) and are motivated to imitate others' actions (Meltzoff & Moore, 1977). Social cues, such as shared attention, support learning about objects and language (Baldwin, 1995; Scaife & Bruner, 1975; Tomasello & Farrar, 1986; Yu & Smith, 2013). In the second-half year of life, infants begin to use social referencing to regulate their own behavior and affect (Feinman, 1982; Gunnar & Stone, 1984; Sorce et al., 1985). Infants are also selective when using others to guide learning (Poulin-Dubois & Brosseau-Liard, 2016). For example, infants selectively interact with new toys when they are presented by an expert, relative to a nonexpert informant (Stenberg, 2013). Other work shows 8-month-old infants are able to track the reliability of an adult's gaze to make predictions about future events (Tummeltshammer et al., 2014).

Social information also influences rule learning abilities in infancy and subsequent executive functions' abilities in later childhood. When 8-month-olds are presented with auditory or visual stimuli that follow an ABA or AAB pattern, they show better extraction of the abstract pattern when presented with speech sounds relative to musical tones or visual shapes (Ferguson & Lew-Williams, 2016; Marcus et al., 2007). Similarly, 7-month-olds show superior abstraction of statistical regularities in syllable sequences when presented with infant-directed versus adult-direct speech (Thiessen et al.,

2005). Other work indicates that 9–10-month-olds use experimenters to guide search behavior in the A-Not-B task (Topal et al., 2008; Werchan & Amso, 2020). More broadly, socioemotional features of infants' environments, such as parenting practices, are predictive of executive functions' ability in early childhood (Blair et al., 2014; Fay-Stammach et al., 2014; Hammond et al., 2012; Rosen et al., 2019). Taken together, these findings indicate that social information may have a unique role in supporting early learning and rule-guided behavior.

Given these findings, we hypothesize that infants may be biased to use social stimuli to structure inputs into hierarchical rule sets. In this view, social information may be a foundational cue that infants exploit to structure learning and help make sense of the abundance of new and often cluttered and ambiguous inputs. In turn, this may scaffold the development of more complex executive functions over ontogenetic development.

To examine whether infants are biased to use social stimuli to structure reinforcement learning, we adapted a task used to examine incidental structure learning in infants (Werchan et al., 2015). We tested 9-month-old infants given that prior work indicates that infants are capable of structure learning (Werchan et al., 2015, 2016) and show reward responses to social information (Tummeltshammer et al., 2019) by this age. Infants were presented with four cue/target-location pairings, in which the cues consisted of static faces presented on colorful circle backgrounds and the target locations consisted of a cartoon animation presented in one of the four screen quadrants (Figure 2). Infants could learn these pairings as individual associations, or they could use the faces or colorful backgrounds as a higher-order context to organize simpler color-location or face-location rules, respectively (Figure 3). We then tested *how* infants learned the pairings in a subsequent generalization task. Given the ostensible value of social stimuli, we predicted that infants would be biased to use the faces as a higher-order context to structure learning. More specific hypotheses and predictions are detailed at the end of the following section.

2 | METHOD

2.1 | Participants

Thirty-nine 9-month-old infants (20 females, 19 males; M age = 9.4 months, SD = 0.46 months) were recruited via advertisements and through birth records from the state department of health. Infants were randomly assigned to a social-context condition (N = 19, 12 females, 7 males; M age = 9.4 months, SD = 0.46 months, 13 non-Hispanic white, 1 black, 2 Hispanic, 3 mixed-race) or a color-context condition (N = 20, 8 females, 12 males; M age = 9.4 months, SD = 0.5 months, 13 non-Hispanic white, 2 black, 5 Hispanic). The sample size was determined using an a priori power analysis at 90% power with a predicted medium effect size (f = 0.25) estimated from prior studies using similar paradigms (Werchan et al., 2015, 2016), which indicated that approximately 40 infants (20 per condition) would provide sufficient power. The final data point was not collected because of the onset of the COVID-19 pandemic. The study was conducted according to guidelines laid down in the Declaration of Helsinki, with written informed consent obtained from a parent or guardian prior to any data collection. All procedures involving human subjects were approved by the Institutional Review Board at Brown University. All infants were born full term (within 4 weeks of due date) and had no history of serious health problems. One additional infant was tested, but their data were discarded because of fussiness and crying. All families were compensated for time and travel.

2.2 | Materials

2.2.1 | Apparatus

Infants sat on a parent's lap approximately 60 cm away from a 24" monitor in a dimly lit room. Stimuli were presented via the SMI Experiment Center software, and infants' eye movements were

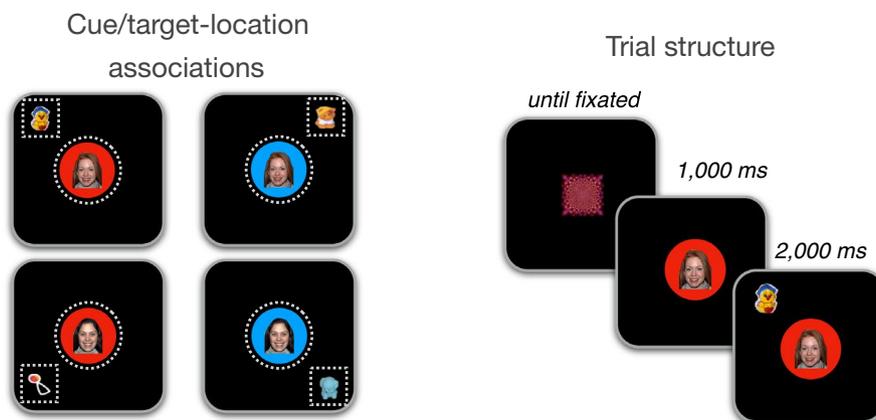


FIGURE 2 Examples of four cue/target-location pairings during the initial learning task (a). The central cues consisted of female faces on colorful circle backgrounds and the target locations consisted of a cartoon animation played in one of four screen quadrants. Eye movement reaction times from the central cue to the target locations were measured (AOIs are illustrated by the dotted gray lines). Each trial began with a colorful attention getter, after which the central cue appeared followed by the animated toy (b).

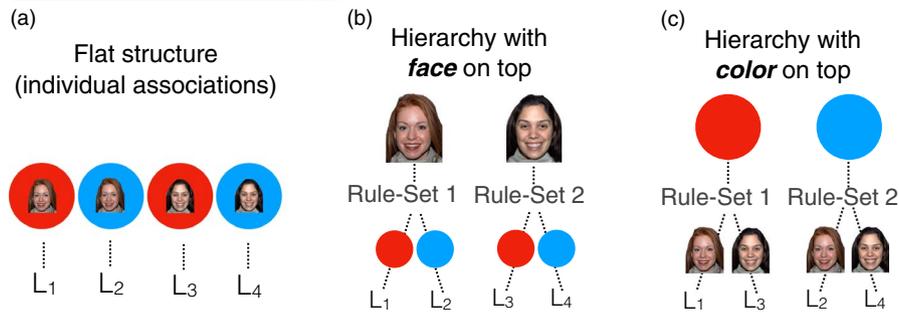


FIGURE 3 Infants could learn the cue/target-location pairings as four separate individual associations (a). Alternatively, infants could apply a hierarchical structure to learn the pairings, using either the faces (b) or the colorful backgrounds (c) as a higher-order context to structure simpler color-location or facelocation rules, respectively. Note that L_1 , L_2 , L_3 , and L_4 refer to the target locations (i.e., screen quadrants in which the cartoon reward appeared).

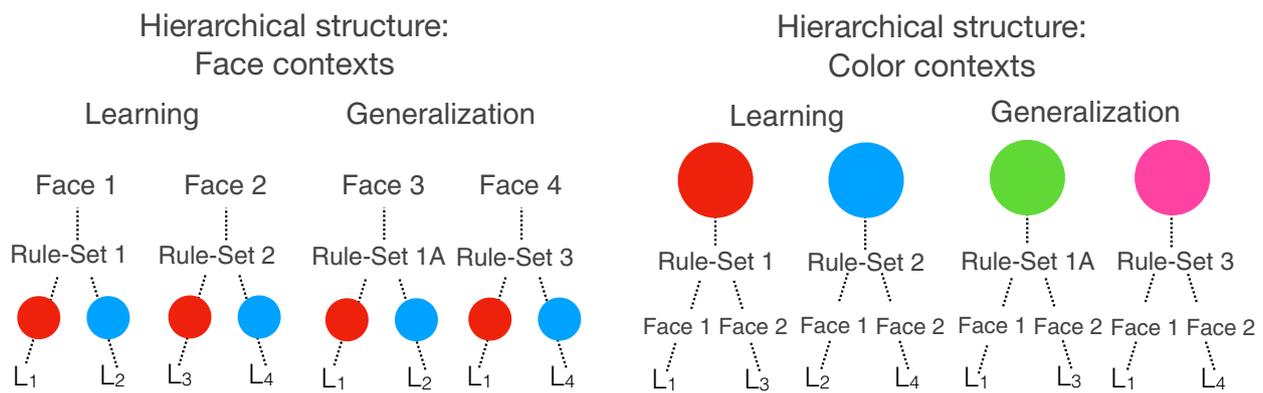


FIGURE 4 Schematic of the hierarchical structure that could be applied during the learning and generalization task using either faces (left panel) or the colored backgrounds (right panel) as a higher-order context. Note that L_1 , L_2 , L_3 , and L_4 refer to the target locations (i.e., screen quadrants in which the cartoon reward appeared).

recorded at a rate of 60 Hz using remote eye tracking software (RED system; SensoMotoric Instruments, or SMI; Teltow, Germany). Prior to the task, each infant's point of gaze (POG) was calibrated by presenting a looming stimulus in the center and each of the four corners of the screen until it was fixated by the infant. Calibration was validated by presenting the same stimulus in four locations on the screen. Infant's estimated POG was compared with the stimulus location, and calibration was repeated if deviations were greater than 2° . Areas of interest (AOIs) were defined in the native SMI software-analysis package BeGaze (Figure 2).

2.2.2 | Stimuli

The stimuli consisted of central cues ($\sim 6.5^\circ$ in size) and targets ($\sim 3.5^\circ$), which were presented on a black screen (Figure 2). The central cues consisted of four discriminable Caucasian female faces (taken from the NimStim Face Stimulus Set, Tottenham et al., 2009) that were overlaid on discriminable colorful circle backgrounds (red, green, blue, pink). The targets consisted of four animated toys (drawn from the Tobii eye tracking calibration stimuli) presented in one of the four screen quadrants, each of which made a unique, discriminable

sound (high-pitched ding, electronic crescendo, low-pitched ringing, horn crescendo), and bounced from side to side. The targets were all equidistant ($\sim 7^\circ$) from the center.

2.3 | Procedure

2.3.1 | Task overview

Infants participated in a learning task followed by a generalization task. All infants received the same learning task, but infants were randomly assigned to either a face-context ($n = 19$) or a color-context ($n = 20$) condition in the generalization task (Figure 4). During the learning task, infants were presented with cue/target-location pairings. The central cues consisted of static female faces overlaid on colorful circle backgrounds, and the targets consisted of rewarding cartoon animations presented in one of the four different screen quadrants. These pairings could be learned simply as individual associations between the central cues and the target locations. Alternatively, infants could apply a hierarchical structure to learn the pairings, using either the faces or the colorful backgrounds as a higher-order context to organize simpler color-location or



face-location rules, respectively. The subsequent generalization task was designed to probe how infants structured the inputs during the initial learning task, namely whether infants learned: (a) individual cue/target associations, (b) a hierarchical structure using the faces to organize simpler color/target-location rules, or (c) a hierarchical structure using the colorful backgrounds to organize simpler face/target-location rules.

The mappings between rule sets, faces, colors, and target locations were counterbalanced across infants. Our measure of interest was infants' average reaction time (RT) to the location of the target (animated toy), which was defined as the time between the target onset and the time infants' point of gaze (i.e., time to first fixation) arrived at the target location. All statistical analyses were carried out using SPSS Version 22.0.

2.3.2 | Learning task

We presented infants with four cue/target-location pairings, in which the centrally presented cues consisted of one of two different female faces presented on one of two colorful circle backgrounds. The four central cues were paired with an animated toy presented in one of the screen quadrants, each of which made a unique sound and bounced. The learning task was designed such that the cue/target-location pairings could be learned in multiple ways. Most simply, infants could learn the pairings as four separate face/color/target-location associations with no latent structure (Figure 3a). This would result in efficient initial learning of the pairings but would not support transfer of prior learning to new stimuli. Alternatively, infants could apply a latent hierarchical structure to learn the pairings, using one of the two dimensions of the central cues (the faces or background colors) as a higher-order context. For instance, infants could use the faces to organize simpler color/target-location rules (Figure 3b). Alternatively, infants could use the colorful backgrounds to organize simpler face/target-location rules (Figure 3c). Although applying a latent hierarchical structure might result in more effortful initial learning, it affords future generalization opportunities, thus allowing for efficient learning in novel contexts.

Infants received 8 trials of the four different cue/target-location pairings for a total of 32 trials (16 in rule set 1 and 16 in rule set 2). Prior to each trial, infants' point of gaze was centered by presenting a colorful central attention-getting stimulus (a looming square fractal making a "boing" sound). The trial was initiated once the experimenter judged that the infant was looking at the central attention-getter. During each trial, the central cue was shown for 1000 ms, after which the animated toy appeared for 2000 ms in the target location associated with the cue (Figure 2). The central cue remained on the screen, while the animated toy was presented. Trial order was pseudo-randomized with the constraint that there were an equal number of trials in which the face changed from one trial to the next and trials in which the colorful background changed from one trial to the next.

Dependent measures

Infants' RTs to the target location from the target onset were measured during the learning task. Trials with eye movement reaction times slower or faster than two standard deviations from the mean for each infant were excluded from analysis. Trials in which infants did not fixate on the central cue prior to attending the target location were also excluded. To smooth over trial-by-trial variability and to account for missing trial data, we binned every two consecutive trials per rule set to create four learning bins for each rule set. Infants contributed an average of 13.56 ($SD = 2.44$) (of 16 possible) trials in rule set 1 and 13.13 ($SD = 2.18$) (of 16 possible) trials in rule set 2.

2.3.3 | Generalization task

Immediately after the learning task, infants were randomly assigned to a face-context generalization condition ($n = 19$) or a color-context generalization condition ($n = 20$), during which infants saw four new cue/target-location pairings (Figure 4). Infants again received 8 pseudorandomized trials of each cue/target-location pairing for a total of 32 trials. The same timing parameters were used in the generalization task as in the learning task.

In the face-context condition, the new central cues consisted of two novel faces presented on the same colorful backgrounds as in the initial learning task. Infants could again use the novel faces to organize the inputs into simpler color/target-location rules (Figure 4). Critically, one novel face denoted a rule set (RS1-A) that had the same previously learned color/target-location rules as a rule set from the initial learning task (RS1). In contrast, the other novel face denoted a novel rule set (RS3) that had two color/target-location rules that had both been experienced individually before, but across different rule sets (RS1 and RS2).

The color-context condition had an analogous design, but the novel cues instead consisted of two novel colorful backgrounds that were paired with the same faces used in the initial learning task (Figure 4). One of these novel colorful backgrounds denoted a rule set (RS1-A) with the same face/target-location rules as a rule set from the learning task. The other colorful background denoted a novel rule set (RS3) that had two face/target-location rules that had been experienced across different rule sets (RS1 and RS2).

Dependent measures

Infants' RTs to the target locations from the target onset were measured. We again binned every two consecutive trials to create four learning bins per rule set. Trials without a fixation to the central cue prior to fixating on the target location or trials with eye movement reaction times slower or faster than two standard deviations from the mean were excluded from analysis. Infants in the face-context condition contributed an average of 12.37 ($SD = 2.59$) (of 16 possible) trials in rule set 1A and 12.21 ($SD = 2.64$) (of 16 possible) trials in rule set 3, and infants in the color-context condition contributed an average of 12.70 ($SD = 3.08$) trials in rule set 1A and 12.05 ($SD = 3.05$) trials in rule set 3.

2.3.4 | Specific predictions

Our predictions were as follows. First, we predicted that infants' RTs should decrease with trial exposure (Canfield & Haith, 1991). Moreover, if infants simply learned individual associations between the central cues and target locations during the learning task, then we expected to observe similar learning rates (i.e., reductions in RTs) between the two novel face or color contexts in the generalization task and no differences by condition assignment. However, if infants applied a hierarchical structure to learn the pairings using the *faces* as a higher-order context, then we expected to find the transfer of prior learning in the face-context condition only (evidenced by faster RTs in RS1A relative to RS3) and no learning or transfer effects in the color-context condition. Additionally, we predicted to find a RT switch cost if infants were learning a hierarchical structure. That is, if infants used the faces as a higher-order context, then we expected to find slower average RTs on trials when the higher-order face rule changes from trial-to-trial relative to average RTs on trials when it remained the same. This prediction is based on prior work, indicating that adults have slower reaction times when a higher-order rule changes on a trial-by-trial basis and thus has to be updated into working memory (Collins & Frank, 2013; Monsell, 2003) and infant work showing the same pattern in a hierarchical rule-learning task structured like the one used here (Werchan et al., 2015). Finally, if infants applied a hierarchical structure using the *colors* as a higher-order context, then we expected to find the transfer of prior learning in the color-context condition only and no learning or transfer effects in the face-context condition.

3 | RESULTS

3.1 | Learning task

We first conducted a mixed-effects ANOVA on infants' averaged reaction times using the rule set (RS1, RS2) and trial bin (bin 1, bin 2, bin 3, bin 4) as within-subjects' factors and condition (face-context, color-context) as a between-subjects factor. This analysis revealed a significant main effect of trial bin, $F(3, 111) = 7.38, p < 0.001, \eta_p^2 = 0.17$, indicating that infants' reaction times became faster with trial exposure (Figure 5). There were no other significant effects or interactions, all $ps > 0.13$, all $Fs < 1.91$, all $\eta_p^2 < 0.05$. These results indicate that infants learned the rule sets with trial exposure and that there were no differences by rule set or condition assignment.

We next examined reaction time costs during face and color change trials in the learning task. We calculated two switch-cost values: (a) a face switch cost (average RTs on trials where the face changed from the previous trial – average RTs on trials where the face remained the same as the previous trial), which assumes a higher-order structure using the faces as a context; (b) a color switch cost (average RTs on trials where the color changed from the previous trial – average RTs on trials where the color remained the same as the previous trial), which assumes a higher-order structure using

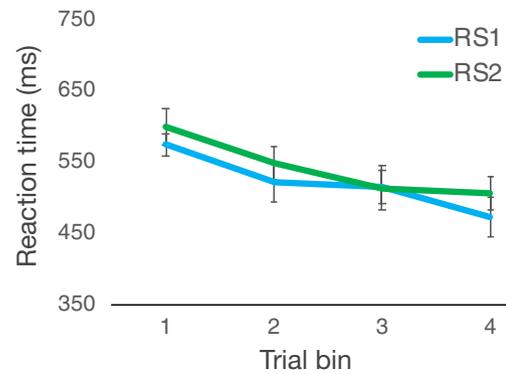


FIGURE 5 Infants' eye movement reaction times over trial bins for RS1 and RS2 during the learning task. Error bars reflect SEM.

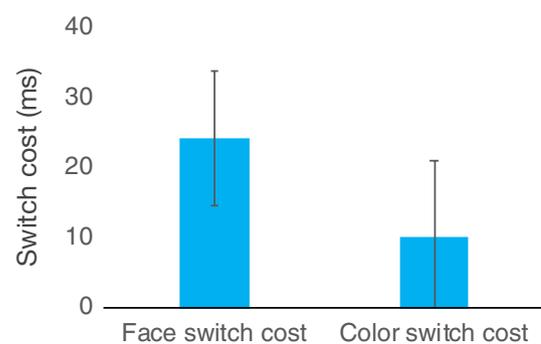


FIGURE 6 Infants' face and color switch costs (average RTs on trials where the face or color, respectively, changes from the previous trial – average RTs on trials where it remains the same). Infants' face switch costs were significantly greater than zero, indicating that their RTs were significantly slower when the face rule changed relative to when it remained the same. Color switch costs did not differ from zero, indicating that there were no differences in infants' RTs when the color changed from trial to trial relative to when it remained the same. Error bars reflect SEM.

color as a context. We first examined whether there were differences in infants' face switch costs relative to color switch costs during learning using a mixed-effects ANOVA with the condition (face-context, color-context) as a between-subjects variable and switch cost type (face switch cost, color switch cost) as a within-subjects' variable. This indicated no significant difference between face and color switch costs, $F(1, 37) = 0.62, p = 0.44, \eta_p^2 = 0.02$. There was also no main effect of condition, $F(1, 37) = 2.10, p = 0.16, \eta_p^2 = 0.05$, nor an interaction between the switch cost type and condition, $F(1, 37) = 0.06, p = 0.82, \eta_p^2 = 0.001$. However, planned comparisons to zero indicated that infants, as a group, had face switch costs significantly greater than zero, $t(38) = 2.55, p = 0.01$, but not for color switch costs, $t(38) = 0.92, p = 0.36$ (Figure 6). This indicates that infants had significantly slower reaction times on trials where the higher-order face changed from one trial to the next relative to when it remained the same, but that there were no differences in reaction times for trials where the color changed relative to trials where the color remained the same. The significant reaction time switch cost

observed on trials when the face changes, requiring infants to update a new rule set into working memory, provides further evidence that infants were learning a hierarchical structure using the faces as a higher-order context.

3.2 | Generalization task

Finally, we examined infants' performance during the generalization task using a mixed-effects ANOVA with the rule set (RS1A, RS3) and trial bin (bin 1, bin 2, bin 3, bin 4) as within-subjects' factors and condition (face-context, color-context) as a between-subjects factor. This analysis revealed a significant main effect of rule set, $F(1, 37) = 8.62$, $p = 0.006$, $\eta_p^2 = 0.20$. Importantly, it also revealed a significant interaction between the rule set and condition, $F(1, 37) = 7.77$, $p = 0.008$, $\eta_p^2 = 0.17$, and a trending interaction between the rule set, trial bin, and condition, $F(3, 111) = 2.46$, $p = 0.07$, $\eta_p^2 = 0.06$. We followed up on this analysis by examining each condition separately (Bonferroni corrected $\alpha = 0.025$). In the face-context condition, a repeated-measures ANOVA revealed a significant main effect of rule set, $F(1, 18) = 26.67$, $p < 0.001$, $\eta_p^2 = 0.60$. There was no significant effect of trial bin, $F(3, 54) = 2.59$, $p = 0.06$, $\eta_p^2 = 0.13$, nor was there a significant interaction between the rule set and trial bin, $F(3, 54) = 1.44$, $p = 0.24$, $\eta_p^2 = 0.07$. In contrast, the color-context condition showed no main effect of rule set, $F(1, 17) = 0.01$, $p = 0.93$, $\eta_p^2 = 0$, or trial bin, $F(3, 57) = 0.47$, $p = 0.71$, $\eta_p^2 = 0.02$. There was also no significant rule set by trial bin interaction, $F(3, 57) = 2.10$, $p = 0.11$, $\eta_p^2 = 0.10$. In sum, these analyses indicated that infants had faster overall reaction times in RS1A relative to RS3 in the face-context generalization condition, but that there were no differences between RS1A and RS3 for infants assigned to the color-context generalization condition (Figure 7). This indicates the transfer of prior learning in the face-context generalization condition, but not in the color-context generalization condition.

4 | DISCUSSION

The ability to structure learning into hierarchically organized rule structures is essential for adaptive behavior. Yet, it is unclear what cues are most likely to be used to structure learning in infancy. In the

current study, we examined whether infants are biased to use social stimuli (in this case faces) to organize inputs into hierarchically organized structures. We first presented infants with cue/target-location pairings during an initial learning task. We found that all infants showed a reduction in reaction times with trial exposure, suggesting that they learned the pairings. We then examined infants' performance in the generalization task to explore *how* infants learned these pairings. Critically, we found evidence that infants structured the inputs into hierarchical rule sets using the faces as a higher-order context. This is shown by the finding that infants in the face-context generalization condition showed the positive transfer of an analogous rule set (as shown by decreased reaction times to RS1A) and negative transfer of a novel rule set (as shown by increased reaction times to RS3). In contrast, infants in the color-context generalization condition showed increased reaction times to both RS1A and RS3 with no differences in performance between rule sets, indicating no transfer of prior learning. Moreover, these results control for the possibility that the decline in reaction times over the learning task was related to a general eye movement speed decline. If this were the case, then we would not expect to observe increased reaction times in the color-context generalization condition as well as in the novel rule set in the face-context generalization condition. As such, these findings indicate that, given the option to use either faces or colorful shapes to structure learning, infants structured learning into hierarchical rule sets using the faces as a higher-order context. In turn, this supported generalization to novel face contexts.

These findings show that infants are biased to use social stimuli to organize reinforcement learning *when in competition with color information*. Other work has found that when non-social stimuli are used, infants can learn to use color or shape information as a higher-order context to structure inputs (Werchan & Amso, 2020; Werchan et al., 2015). It is likely that infants can learn to use different information to structure learning as a function of changing task demands. For instance, this paradigm could be adapted to test whether infants are biased to use shape as a higher-order context if a label is heard during learning.

Interestingly, while we found that infants in the face-context generalization condition showed positive transfer (faster reaction times to RS1A relative to RS3), we also observed that infants' reaction times did not show a further decrease over trials during the generalization task. The lack of a further increase in reaction

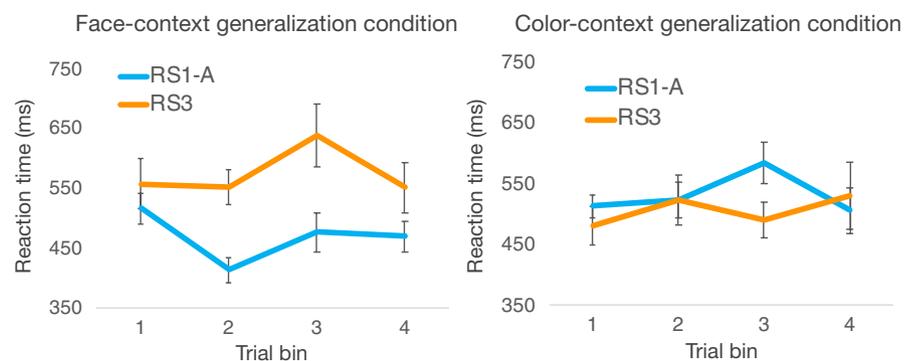


FIGURE 7 Infants' eye movement reaction times over trial bins for RS1-A and RS3 during the generalization task for the face-context generalization condition and the color-context generalization condition. Error bars reflect SEM.

times could suggest that this mechanism is not sufficiently mature to support proactive or anticipatory saccades. This explanation is consistent with computational approaches, which suggest that proactive control over action requires more complex hierarchical nesting of frontostriatal circuitry (Collins & Frank, 2013). It is also possible that the 1000 ms delay between the cue and target onset was too short to elicit anticipatory eye movements or that the cue remaining on screen during the delay decreased the number of anticipatory eye movements due to the added cost of disengaging from a salient cue. Alternatively, the use of faces could have also altered infants' attentional patterns, potentially leading to asymptotic reaction times. For instance, prior work suggests that the use of salient faces in spatial cueing tasks can attenuate typical inhibition-of-return effects in both adult (Pérez-Dueñas et al., 2014) and infant studies (Markant et al., 2015). Future work is needed to address these questions.

We also measured infants' reaction time switch costs during face and color change trials in the learning task based on prior adult (Monsell, 2003) and infant (Werchan et al., 2015) studies showing that individuals have slower reaction times when updating a new higher-order rule into working memory. We found that infants' face switch costs were significantly greater than zero. As an analytical control, we also examined color switch costs and saw no differences in reaction times. These data are consistent with infants' generalization performance and provide further evidence that infants used the faces to structure inputs into hierarchical rule sets.

Prior research argues that humans are evolutionarily adapted to transmit knowledge through pedagogical interactions with caregivers and other conspecifics (Csibra & Gergely, 2006). This idea is supported by findings showing that human children use social cues to guide attention and learning from early infancy. For instance, social cues such as joint attention have been shown to guide object processing (Cleveland et al., 2007; Reid & Striano, 2005; Striano et al., 2006; Theuring et al., 2007; Wu et al., 2011), multimodal learning (Wu & Kirkham, 2010), and event sequence learning (Topal et al., 2008; Werchan & Amso, 2020a). We add to these findings by showing that 9-month-old infants can use static faces to organize arbitrary cue/target pairings into hierarchical rule sets.

Establishing the value of social information for structuring reinforcement learning beginning in early postnatal life also has implications for understanding the ontogenetic origins of more complex behavior and executive functions. For example, prior work indicates that young children monitor the reliability of informants to infer whether to trust individuals (Clement et al., 2004; Koenig & Harris, 2005; Luchkina et al., 2018; Tummeltshammer et al., 2014; Xiao et al., 2018). Our findings suggest that this reinforcement learning mechanism could scaffold learning of more abstract rules, such as rules related to the trustworthiness. This adds insight into the impact of early social environments on adaptive functioning more broadly. For example, if the quality of interactions with caregivers is poor or unreliable, this may alter infants' bias to use social contexts to guide learning as an adaptive response to early caregiving environments, as predicted by ecological approaches to executive functions'

development (Werchan & Amso, 2017). In turn, this may influence children's capacity to engage with and use other children and adults in the learning process, potentially impacting academic achievement as the child begins formal schooling. This idea is supported by work indicating predictive relationships between the quality of early parent-child interactions on subsequent school readiness, which is mediated by the development of executive functions (Blair & Raver, 2015; Devine et al., 2016; Fay-Stammbach et al., 2014; Russell et al., 2016). Moreover, if infants use social information to structure learning, then it is possible that infants might use caregivers and related social stimuli to learn more complex rule sets, such as those guiding social interactions, cultural norms, and even parenting practices. This could lead to phenomena such as "ghosts in the nursery," where individuals model caregiving practices after those they received during early childhood (Fraiberg et al., 1975). The present work adds new insights into the potential mechanistic origins of these observed relationships.

A limitation of the present study is that we are unable to tease apart the precise mechanism underlying infants' bias to use social stimuli to structure learning. Social stimuli are both attention-capturing (Amso et al., 2014; Frank et al., 2009; Gluckman & Johnson, 2013; Tummeltshammer, Wu, et al., 2014) and are associated with the reward value for learning (Tummeltshammer et al., 2019). Additionally, it is possible that the figure-ground design of the face/color cues could have caused infants to bias attention to the face during learning. More work is needed to discern these possibilities. For instance, future research could directly measure individual differences in infants' baseline attentional biases and physiological reward responses to both social and non-social stimuli and examine whether these individual differences predict infants' bias to use faces to structure learning. Nonetheless, the current findings provide an important advance in understanding how infants use simultaneously presented social and non-social information to structure learning of hierarchical rule structures.

In sum, here we found that infants are biased to use social stimuli to organize inputs into hierarchical rule sets that support generalization to new contexts. These findings provide new evidence that social information helps us to narrow the space of learning problems by acting as a relevant cue that infants leverage to organize inputs into flexible rule structures. While these findings add important new insights into the role of social information in structuring learning in infancy, the generalizability of our findings is limited by our sample demographic, which was predominately white and middle class. We designed the study to use social stimuli that matched the race and gender of the majority of primary caregivers in our infant sample, which primarily consisted of white female caregivers. We made this design choice based on prior work showing that mothers elicit a reward response (Tummeltshammer et al., 2019). However, given that our infant sample is not representative of the population at large, this limits the generalizability of our findings to non-white middle class infants. Future work is needed with more representative samples to determine the generalizability and relevance of the current findings for the population at large.



ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health (R21 MH113870 to DA) and NSF (2051819 to DA). We thank Jamie Poldrack and Lily Gordon for help with recruitment and data processing and all of the infants and families who made this research possible.

CONFLICTS OF INTEREST

The authors report no conflicts of interest.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon reasonable request.

ORCID

Denise M. Werchan  <https://orcid.org/0000-0001-8393-5933>

REFERENCES

- Amso, D., Haas, S., & Markant, J. (2014). An eye tracking investigation of developmental change in bottom-up attention orienting to faces in cluttered natural scenes. *PLoS ONE*, *9*(1), 1–7. <https://doi.org/10.1371/journal.pone.0085701>
- Baldwin, D. A. (1995). Understanding the link between joint attention and language. In C. Moore & P. J. Dunham (Eds.), *Joint attention: Its origins and role in development* (pp. 131–158). Erlbaum.
- Blair, C., & Raver, C. C. (2015). School readiness and self-regulation: A psychobiological approach. *Annual Review of Psychology*, *66*, 711–731. <https://doi.org/10.1146/annurev-psych-010814-015221>
- Blair, C., Raver, C. C., & Berry, D. J. (2014). Two approaches to estimating the effect of parenting on the development of executive function in early childhood. *Developmental Psychology*, *50*(2), 554–565. <https://doi.org/10.1037/a0033647>
- Canfield, R. L., & Haith, M. M. (1991). Young infants' visual expectations for symmetric and asymmetric stimulus sequences. *Developmental Psychology*, *27*(2), 198–208.
- Clement, F., Koenig, M., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language*, *19*(4), 360–379.
- Cleveland, A., Schug, M., & Striano, T. (2007). Joint attention and object learning in 5- and 7-month-old infants. *Infant and Child Development*, *16*, 295–306. <https://doi.org/10.1002/icd>
- Collins, A. G. E., Cavanagh, J. F., & Frank, M. J. (2014). Human EEG uncovers latent generalizable rule structure during learning. *Journal of Neuroscience*, *34*(13), 4677–4685. <https://doi.org/10.1523/JNEUROSCI.3900-13.2014>
- Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229. <https://doi.org/10.1037/a0030852>
- Cooper, R. P., & Aslin, R. N. (1990). Preference for infant-directed speech in the first month after birth. *Child Development*, *61*(5), 1584–1595.
- Csibra, G., & Gergely, G. (2006). Social learning and social cognition: The case for pedagogy. In Y. Munakata & M. H. Johnson (Eds.), *Processes of change in brain and cognitive development* (pp. 249–274). Oxford University Press.
- Devine, R. T., Bignardi, G., & Hughes, C. (2016). Executive function mediates the relations between parental behaviors and children's early academic ability. *Frontiers in Psychology*, *7*, 1–15. <https://doi.org/10.3389/fpsyg.2016.01902>
- Farroni, T., Massaccesi, S., Pividori, D., Simion, F., & Johnson, M. H. (2004). Gaze following in newborns. *Infancy*, *5*(1), 39–60.
- Fay-Stammbach, T., Hawes, D. J., & Meredith, P. (2014). Parenting influences on executive function in early childhood: A review. *Child Development Perspectives*, *8*(4), 258–264. <https://doi.org/10.1111/cdep.12095>
- Feinman, S. (1982). Social referencing in infancy. *Merrill-Palmer Quarterly*, *28*(4), 445–470.
- Ferguson, B., & Lew-Williams, C. (2016). Communicative signals support abstract rule learning by 7-month-old infants. *Nature Scientific Reports*, *6*, 25434. <https://doi.org/10.1038/srep25434>
- Fraiberg, S., Adelson, E., & Shapiro, V. (1975). Ghosts in the nursery: A psychoanalytic approach to the problems of impaired infant-mother relationships. *Journal of American Academy of Child Psychiatry*, *14*(3), 387–421.
- Frank, M. C., Vul, E., & Johnson, S. P. (2009). Development of infants' attention to faces during the first year. *Cognition*, *110*(2), 160–170. <https://doi.org/10.1016/j.cognition.2008.11.010>
- Gluckman, M., & Johnson, S. P. (2013). Attentional capture by social stimuli in young infants. *Frontiers in Psychology*, *4*, 1–7. <https://doi.org/10.3389/fpsyg.2013.00527>
- Gunnar, M. R., & Stone, C. (1984). The effects of positive maternal affect on infant responses to pleasant, ambiguous, and fear-provoking toys. *Child Development*, *55*(4), 1231–1236.
- Hammond, S. I., Müller, U., Carpendale, J. I. M., Bibok, M. B., & Liebermann-Finestone, D. P. (2012). The effects of parental scaffolding on preschoolers' executive function. *Developmental Psychology*, *48*(1), 271–281. <https://doi.org/10.1037/a0025519>
- Hood, B. M., Willen, J. D., & Driver, J. (1998). Adult's eyes trigger shifts of visual attention in human infants. *Psychological Science*, *9*(2), 131–134.
- Johnson, M. H. (2005). Subcortical face processing. *Nature Reviews Neuroscience*, *6*, 766–774. <https://doi.org/10.1038/nrn1766>
- Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277.
- Luchkina, E., Sobel, D. M., & Morgan, J. L. (2018). Eighteen-month-olds selectively generalize words from accurate speakers to novel contexts. *Developmental Science*, *21*, e12663. <https://doi.org/10.1111/desc.12663>
- Marcus, G. F., Fernandes, K. J., & Johnson, S. P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, *18*(5), 387–391.
- Markant, J., Oakes, L. M., & Amso, D. (2015). Visual selective attention biases contribute to the other-race effect among 9-month-old infants. *Developmental Psychobiology*, *58*(3), 355–365. <https://doi.org/10.1002/dev.21375>
- Meltzoff, A. N., Kuhl, P. K., Movellan, J., & Sejnowski, T. J. (2009). Foundations for a new science of learning. *Science*, *325*, 284–289.
- Meltzoff, A. N., & Moore, K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, *198*(4312), 75–78.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. [https://doi.org/10.1016/S1364-6613\(03\)00028-7](https://doi.org/10.1016/S1364-6613(03)00028-7)
- Pérez-Dueñas, C., Acosta, A., & Lupiáñez, J. (2014). Reduced habituation to angry faces: Increased attentional capture as to override inhibition of return. *Psychological Research Psychologische Forschung*, *78*, 196–208. <https://doi.org/10.1007/s00426-013-0493-9>
- Poulin-Dubois, D., & Brosseau-Liard, P. (2016). The developmental origins of selective social learning. *Current Directions in Psychological Science*, *25*(1), 60–64. <https://doi.org/10.1177/0963721415613962>
- Reid, V. M., & Striano, T. (2005). Adult gaze influences infant attention and object processing: Implications for cognitive neuroscience. *European Journal of Neuroscience*, *21*, 1763–1766. <https://doi.org/10.1111/j.1460-9568.2005.03986.x>
- Rosen, M. L., Amso, D., & Mclaughlin, K. A. (2019). The role of the visual association cortex in scaffolding prefrontal cortex development: A novel mechanism linking socioeconomic status and executive

- function. *Developmental Cognitive Neuroscience*, 39(July), 100699. <https://doi.org/10.1016/j.dcn.2019.100699>
- Russell, B. S., Lee, J. O., Spieker, S., & Oxford, M. L. (2016). Parenting and preschool self-regulation as predictors of social emotional competence in 1st grade. *Journal of Research in Childhood Education*, 30(2), 153–169. <https://doi.org/10.1080/02568543.2016.1143414>
- Scaife, M., & Bruner, J. S. (1975). The capacity for joint visual attention in the infant. *Nature*, 253(5489), 265–266.
- Sorce, J. F., Emde, R. N., Campos, J., & Klinnert, M. D. (1985). Maternal emotional signaling: Its effect on the visual cliff behavior of 1-year-olds. *Developmental Psychology*, 21(1), 195–200.
- Stenberg, G. (2013). Do 12-month-old infants trust a competent adult? *Infancy*, 18(5), 873–904. <https://doi.org/10.1111/infa.12011>
- Striano, T., Chen, X., Cleveland, A., & Bradshaw, S. (2006). Joint attention social cues influence infant learning. *European Journal of Developmental Psychology*, 3(3), 289–299. <https://doi.org/10.1080/17405620600879779>
- Theuring, C., Gredeba, G., & Hauf, P. (2007). Object processing during a joint gaze following task. *European Journal of Developmental Psychology*, 4(1), 65–79. <https://doi.org/10.1080/17405620601051246>
- Thiessen, E. D., Hill, E. A., & Saffran, J. R. (2005). Infant-directed speech facilitates word segmentation. *Infancy*, 7(1), 53–71.
- Tomasello, M., & Farrar, M. J. (1986). Joint attention and early language author. *Child Development*, 57(6), 1454–1463.
- Topal, J., Gergely, G., Miklosi, A., Erdohegyi, A., & Csibra, G. (2008). Infants' perseverative search errors are induced by pragmatic misinterpretation. *Science*, 321(September), 1831–1835.
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., Marcus, D. J., Westerlund, A., Casey, B. J., & Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168(3), 242–249.
- Tummeltshammer, K. S., Feldman, E. C. H., & Amso, D. (2019). Using pupil dilation, eye-blink rate, and the value of mother to investigate reward learning mechanisms in infancy. *Developmental Cognitive Neuroscience*, 36, <https://doi.org/10.1016/j.dcn.2018.12.006>
- Tummeltshammer, K. S., Mareschal, D., & Kirkham, N. Z. (2014). Infants' selective attention to reliable visual cues in the presence of salient distractors. *Child Development*, 85(5), 1981–1994. <https://doi.org/10.1111/cdev.12239>
- Tummeltshammer, K. S., Wu, R., Sobel, D. M., & Kirkham, N. Z. (2014). Infants track the reliability of potential informants. *Psychological Science*, 25(9), 1730–1738. <https://doi.org/10.1177/0956797614540178>
- Werchan, D. M., & Amso, D. (2017). A novel ecological account of prefrontal cortex functional development PFC: The state of the art. *Psychological Review*, 124(6), 720–739.
- Werchan, D. M., & Amso, D. (2020a). Adaptive rule learning of event sequences during the A-not-B task in 9-month-old infants. *Developmental Psychobiology*, 62(8), 1021–1034. <https://doi.org/10.1002/dev.21999>
- Werchan, D. M., & Amso, D. (2020b). Top-down knowledge rapidly acquired through abstract rule learning biases subsequent visual attention in 9-month-old infants. *Developmental Cognitive Neuroscience*, 42, 100761. <https://doi.org/10.1016/j.dcn.2020.100761>
- Werchan, D. M., Collins, A. G. E., Frank, M. J., & Amso, D. (2015). 8-month-old infants spontaneously learn and generalize hierarchical rules. *Psychological Science*, 26(6), 805–815. <https://doi.org/10.1177/0956797615571442>
- Werchan, D. M., Collins, A. G. E., Frank, M. J., & Amso, D. (2016). Role of prefrontal cortex in learning and generalizing hierarchical rules in 8-month-old infants. *Journal of Neuroscience*, 36(40), 10314–10322. <https://doi.org/10.1523/JNEUROSCI.1351-16.2016>
- Wu, R., Gopnik, A., Richardson, D. C., & Kirkham, N. Z. (2011). Infants learn about objects from statistics and people infants learn about objects from statistics and people. *Developmental Psychology*, 47(5), 1220–1229. <https://doi.org/10.1037/a0024023>
- Wu, R., & Kirkham, N. Z. (2010). No two cues are alike: Depth of learning during infancy is dependent on what orients attention. *Journal of Experimental Child Psychology*, 107, 118–136. <https://doi.org/10.1016/j.jecp.2010.04.014>
- Xiao, N. G., Wu, R., Quinn, P. C., Liu, S., Tummeltshammer, K. S., Kirkham, N. Z., & Lee, K. (2018). Infants rely more on gaze cues from own-race than other-race adults for learning under uncertainty. *Child Development*, 89(3), 1–24. <https://doi.org/10.1111/cdev.12798>
- Yu, C., & Smith, L. B. (2013). Joint attention without gaze following: Human infants and their parents coordinate visual attention to objects through eye-hand coordination. *PLoS ONE*, 8(11), e79659. <https://doi.org/10.1371/journal.pone.0079659>

How to cite this article: Werchan DM, Amso D. All contexts are not created equal: Social stimuli win the competition for organizing reinforcement learning in 9-month-old infants. *Dev Sci*. 2021;00:e13088. <https://doi.org/10.1111/desc.13088>